Temporal Intent Reconstruction for Long-Horizon Agentic Predictive Control

Krishna Teja Areti
Fast Enterprises Llc.
North York, ON, Canada

Vijay Putta
Fast Enterprises Llc.
Louisiana, USA

Prudhvi Ratna Badri Satya
Cloudflare Inc.
Texas, USA

Ajay Guyyala
Meta Platforms Inc.
Texas, USA

ABSTRACT

Temporal Intent Reconstruction framework integrated with a Masked Cognitive Predictor to improve predictive control under changing goals and dynamic conditions. Using real multimodal data from HARMONIC, RoboMind, RoboNet, and Open X-Embodiment, the model reconstructs past intent trajectories and embeds misalignment signals into the control objective for long-horizon adaptation. Experiments showed stable reconstruction across embodiment and modality variations, reduced goal divergence by 31.4%, and improved tracking behaviour by 78% during transitions. The framework improved accuracy, RMSE reduction, and tracking behaviour compared with baseline MPC, inverse learning, and reinforcement-based controllers. These results indicate that temporal intent reconstruction enhances consistency and long-range predictive capability in systems operating under varied sensing, morphology, and task settings.

General Terms

Autonomous systems; Predictive control, Cognitive modeling, Temporal reasoning, Robotics, Machine learning, Human–robot interaction, Adaptive control, Multimodal sensing, Intent modeling

Keywords

Temporal Intent Reconstruction, Predictive Control, Cognitive Modeling, Intent Misalignment, Adaptive Robotics, Multimodal Datasets

1. INTRODUCTION

Autonomous and adaptive systems must operate with stable and purposeful control in dynamic environments [13]. Traditional predictive models optimize state transitions but do not reconstruct how goals shift over time, which limits long-horizon adaptability in uncertain settings [40]. This gap becomes clearer in settings where goals, constraints, and task priorities shift [23]. Reconstructing internal intent paths has been suggested as a way to support long-term adaptability [43], since control systems must respond to external signals while inferring how past objectives shape future goals over extended horizons [10]. Model Predictive Control remains a widely used method because of its structured constraint handling [11], yet its fixed cost terms and explicit physical models limit flexibility in changing contexts [27]. Modern robotic and vehicular systems therefore require awareness of intent rather than only state transitions [22]. Reconstructing temporal intent patterns supports earlier prediction of goal divergence [5] and shifts control from reactive prediction to intent-based reasoning, improving adaptability in unstructured environments [4].

The study is driven by the gap between state-based optimization and cognitively inspired prediction [14]. Traditional models optimize trajectories without explaining why an agent selects particular actions [33], whereas humans rely on memory and retrospective reasoning to anticipate changes in goals [18]. Temporal reconstruction enables controllers to infer latent intentions from past actions and predict future behavioral shifts, supporting long-horizon planning through structured decision memory [28].

The limits of current predictive models appear in uncertain settings [17], where MPC and reinforcement methods focus on visible outcomes rather than reconstructing intent paths [31]. This creates sensitivity to short-term variations and weak long-sequence coherence. A system must recreate latent intent, anticipate goal shifts, and adapt its optimization process [15, 45]. Existing inverse and learning-based MPC approaches infer costs from demonstrations [3, 39], but they assume fixed goals, while hybrid controllers remain short-horizon and reactive [21], lacking mechanisms to reconstruct developing intent.

Recent probabilistic and learning-based predictive models introduce neural estimators, Gaussian Processes, and adaptive control barriers. These methods improve uncertainty estimation but still treat intent as an implicit factor. They adjust weights or uncertainty terms but do not rebuild the internal motivation path that explains sequential decisions. Without this reconstruction, interpretability decreases and early detection of behavioral drift becomes difficult, especially in long-horizon tasks. This highlights the need for control frameworks that integrate temporal intent reconstruction directly into the optimization process.

The paper presents a Temporal Intent Reconstruction framework integrated into an agentic predictive control structure, where a Masked Cognitive Predictor reconstructs intent paths from past control histories and predicts goal divergence in latent cognitive space. The system uses intent-misalignment vectors for long-horizon adaptation, offering stable goal-focused behavior without retraining and improving over classical MPC and reinforcement approaches by introducing reconstructive reasoning into time-dependent control. This work is intended to create a Temporal Intent Reconstruction framework in MCP that will allow agents to recreate past intent trajectories, forecast future goal divergence and adjust control behavior in more dynamic objectives and uncertain environments.

To guide the development and analysis of the framework, the research goals outline the model design, operational structure, and adaptability under changing temporal and agentic conditions.

—To model a Masked Cognitive Predictor capable of reconstructing past intent paths and misalignment vectors in a latent cognitive space.

—To integrate reconstructed intent representations into a predictive control architecture that adapts to varying objectives and environments.

—To examine the effectiveness of temporal intent reconstruction compared with existing predictive control models in adaptability, accuracy, and goal alignment.

These objectives connect directly to research questions supporting evaluation of reconstruction fidelity and adaptive performance.

(1) How can past intent trajectories be reconstructed and projected into a latent cognitive space to predict future goal divergence within adaptive control?
(2) How can temporal intent reconstruction be integrated into predictive control to enable dynamic goal realignment under uncertainty?
(3) To what extent does temporal intent reconstruction improve control performance, adaptability, and long-horizon stability compared with standard predictive models?

This research integrates temporal intent reconstruction into predictive control, allowing systems to anticipate goal changes rather than react to state deviations. The framework aligns optimization with interpretive reasoning by reconstructing past cognitive traces, supporting long-term stability without reconfiguration or retraining. This view treats adaptability as an internal property derived from temporal reconstruction, improving transparency, continuity, and context-aware decision making in autonomous control. It strengthens the theoretical basis for adaptive autonomy and supports more robust behaviour in fields such as robotics, process control, and intelligent navigation.

The remainder of this paper is organized as follows. Section 2 reviews prior work on temporal reconstruction, intent modeling, and predictive control. Section 3 details the Temporal Intent Reconstruction framework and its integration with the predictive control architecture. Section 4 describes the datasets, preprocessing steps, and model configuration. Section 5 presents the quantitative evaluation and comparative analysis across domains and embodiments. Section 6 concludes the paper and outlines directions for future research.

2. LITERATURE REVIEW

A transformer-based temporal intent model was developed by [16] using contextual human motion data with a ten-head attention encoder, achieving about 0.18,m and 0.13,m L2 errors and nearly 90% accuracy. [34] applied a GRU–MLP encoder for predicting human approach intentions, reaching an F1 score of 0.75 and a Cohen's κ of 0.835. Visual predictive manipulation was advanced by [19] through action decomposition and a variational autoencoder trained on RoboNet, while [24] introduced the multimodal HARMONIC dataset combining gaze, speech, EEG, and robot-state signals to study shared intent in assistive tasks. Additional temporal modeling efforts included EEG-based fairness analysis by [37], QR–Kalman spatio-temporal reconstruction by [6], and temporal feature learning with FLAN–T5 and LoRA in [32]. The transformer model proposed by [42] also addressed intent prediction but showed overfitting due to limited data diversity.

In the domain of predictive control, [49] highlighted how multi-source temporal reasoning enhances environmental modeling. Inverse and hybrid intent-aware control strategies were expanded by [46], who introduced an inverse MCP using bilevel optimization and the Pontryagin Maximum Principle to estimate Q and R matrices, and by [26], who reviewed data-driven optimal control integrating MCP, reinforcement learning, and hybrid schemes. Bayesian actor–critic control for energy systems appeared in [44], while [12] used inverse reinforcement learning to infer driver intent in autonomous driving. Learning-based MCP extensions included nonlinear system identification by [41], MCP–DRL coordination for power plants by [8], adaptive penalty actor–critic tracking by [36], and subspace-identified MCP for converters in [25]. Gaussian Process–based MCP developments included dynamic-forgetting adaptation for underwater vehicles by [2] and probabilistic uncertainty quantification by [20]. Safety-critical extensions combining Gaussian Processes with Control Barrier Functions were presented by [48], while [35] applied Bayesian multi-task learning for reconstructing structural-health data. Reinforcement-driven predictive control further contributed nonlinear MCP for reactor temperature regulation in [29], linearized MCP for tubular reactors in [47], LSTM–DRL temporal modeling for skid-steer robots in [1], and tube-based MPC enabling precise spacecraft rendezvous in [30].

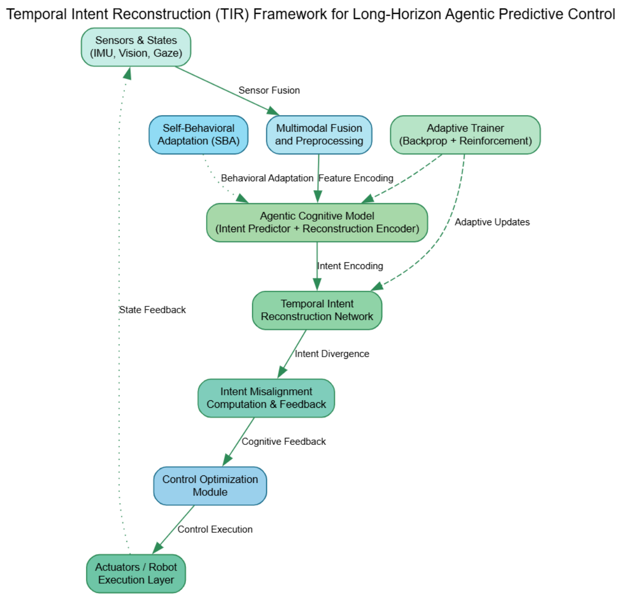

Fig. 1: Architecture of the Temporal Intent Reconstruction framework integrated into the predictive control loop.

3. PROPOSED METHODOLOGY

The methodology introduces the framework for Temporal Intent Reconstruction integrated within MCP, designed to reconstruct latent cognitive trajectories and predict future goal divergence under dynamic environments. The approach builds upon the foundation of traditional predictive control while incorporating intent-based reasoning as an internal feedback layer. It models both the physical and motivational states of an agent, where reconstructed intent vectors are used to refine control optimization across time horizons. This integration enables a system to adapt its decision pathway by interpreting its historical intent, rather than relying solely on immediate error correction. The following subsections describe the mathematical formulation, reconstruction process, predictive optimization, cognitive alignment, and complete adaptive control structure that together establish the proposed methodology.

Fig. 1 illustrates the interaction across perceptual, cognitive, and execution layers in the proposed framework. It shows how temporal intent reconstruction is integrated into the predictive control loop, where perceptual inputs are encoded into latent intent representations, cognitive modules compute intent misalignment signals, and the execution layer adapts control actions to maintain goal-directed and stable behavior over time.

3.1 Problem Formulation

$$x_{t+1} = f(x_t, u_t) + w_t \tag{1}$$

The nonlinear dynamic model describes the temporal development of state variables where x_t denotes the current state, u_t the control input, and w_t the process noise. The transition function $f(\cdot)$ captures how the system responds to control signals under disturbances. This model allows inclusion of latent intent as part of the state transition process, creating an adaptive control foundation. It defines how prior decisions influence future dynamics and sets the stage for intent-based feedback as described in Eq1.

$$J_t = \sum_{k=0}^{N-1} \left[(x_{t+k} - x_{t+k}^{ref})^\top Q (x_{t+k} - x_{t+k}^{ref}) + u_{t+k}^\top R u_{t+k} \right] \tag{2}$$

The control cost J_t quantifies performance by combining state error and control effort across a prediction horizon N. The matrices Q and R represent the penalties on deviation and control energy, respectively. This term balances performance efficiency and control smoothness across time. It serves as the optimization core for predictive control before integrating cognitive feedback. The resulting structure provides a consistent link between physical dynamics and learned intent trajectories (Eq2).

3.2 Temporal Intent Reconstruction

$$\hat{\iota}_t = \Phi(h_{t-1}, x_t, u_t) \tag{3}$$

The reconstructed intent $\hat{\iota}_t$ represents a latent encoding derived from prior hidden states, current states, and applied controls. The function $\Phi(\cdot)$ transforms this temporal context into a vector expressing inferred intent. This captures long-term dependencies across sequential decisions. Through this mapping, the system reconstructs motivation patterns that drive agent behavior. The encoding mechanism provides a bridge between historical reasoning and current objectives, as represented in Eq3.

$$L_{rec} = \|\hat{\iota}_t - \iota_t^{ref}\|_2^2 \tag{4}$$

The reconstruction loss L_{rec} measures the Euclidean distance between predicted and reference intents, encouraging accurate temporal recall. By minimizing this loss, the framework aligns learned intent with its reference target. This consistency helps maintain the agent's contextual understanding over time. The process strengthens temporal coherence during sequential updates. Accurate intent reconstruction stabilizes cognitive feedback within the control framework (Eq4).

3.3 Intent-Guided Predictive Control

$$\delta_t = \hat{\iota}_t - \iota_t^{ref} \tag{5}$$

The misalignment δ_t quantifies deviation between reconstructed and reference intent, representing internal goal drift. It functions as an adaptive correction term that modulates predictive optimization. This component allows the controller to react not only to physical deviations but also to cognitive inconsistencies. It translates reconstructed intent into quantitative guidance for the control layer. The correction process refines long-horizon adaptability through Eq5.

$$J_t^* = \min_{u_t,\dots,u_{t+N}} \sum_{k=0}^{N-1} [\|x_{t+k} - x_{t+k}^{ref}\|_Q^2 + \|u_{t+k}\|_R^2 + \lambda\|\delta_{t+k}\|^2] \tag{6}$$

The extended cost J_t^* incorporates intent divergence within the predictive optimization. The new regularization term $\lambda\|\delta_{t+k}\|^2$ penalizes misalignment while retaining physical accuracy. This multi-objective formulation integrates cognitive and mechanical consistency in one optimization loop. The control decisions adjust dynamically according to reconstructed intent shifts. The hybrid objective described in Eq6 allows intent-aware control refinement.

Algorithm 1 Intent-Guided Predictive Optimization

```
1: Input: Current state x_t, intent î_t, misalignment δ_t, and system model f(·)
2: Output: Optimal control sequence {u_t, ..., u_{t+N}}
3: for prediction horizon N do
4:     Generate initial control guess u_t^0
5:     for iteration i = 1 to I do
6:         Predict next state x_{t+1} = f(x_t, u_t^{i-1})
7:         Compute cost J_t^* using Eq6
8:         Calculate gradient ∇_{u_t} J_t^* and update control u_t^i = u_t^{i-1} − η∇_{u_t} J_t^*
9:         if |J_t^i − J_t^{i-1}| < ε then
10:            Break optimization loop
11:        end if
12:    end for
13:    Apply u_t^* to system and shift horizon
14: end for
15: Return updated control inputs
```

Algorithm 1 minimizes the augmented cost in Eq6 using iterative gradient optimization. It predicts trajectories across horizons while updating control signals. The algorithm allows real-time adaptation through misalignment correction. It provides both predictive efficiency and goal-awareness under dynamic contexts. Each control cycle refines its trajectory to maintain temporal alignment with reconstructed intent.

3.4 Temporal Regularization and Stability

$$L_{temp} = \sum_t \|\hat{\iota}_t - \hat{\iota}_{t-1}\|_2^2 \tag{7}$$

The temporal regularization L_{temp} smooths transitions between consecutive intents to maintain consistent reasoning. It discourages abrupt cognitive fluctuations and supports gradual temporal development. Regularization enhances the stability of reconstructed trajectories. This term prevents noisy oscillations in intent embedding. It reinforces continuity of reconstructed motivation as expressed in Eq7.

$$V(x_{t+1}) - V(x_t) \leq -\alpha \|x_t - x_t^{ref}\|^2 \tag{8}$$

Lyapunov stability in Eq8 confirms closed-loop convergence of the control system. The parameter α defines the rate of decay toward equilibrium. This inequality guarantees that deviations decrease monotonically under intent-based feedback. Stability verification provides theoretical grounding for real-time adaptation. It confirms that reconstructed intent integration preserves system stability under all operational conditions.

3.5 Cognitive Projection and Alignment

$$z_t = W_\iota \hat{\iota}_t + b_\iota \tag{9}$$

Projection into latent cognitive space is achieved through Eq9, where the reconstructed intent is transformed using learnable parameters W_ι and b_ι. This embedding compresses motivational context into a lower-dimensional form. It allows direct comparison between internal and reference cognitive states. The projection enables interpretation of reconstructed intent relative to its task context. It also facilitates alignment between high-level reasoning and control policy representation.

$$L_{align} = \|z_t - z_t^{ref}\|_2^2 + \beta \|\nabla z_t\|_F^2 \tag{10}$$

The alignment loss L_{align} maintains similarity between projected and reference embeddings while penalizing irregular gradients. The second term $\beta\|\nabla z_t\|_F^2$ smooths cognitive transitions. This confirms stable adaptation of cognitive structures across time. The alignment objective connects reconstructed representations with operational consistency. It refines interpretability and coherence in latent reasoning spaces, as formalized in Eq10.

3.6 Optimization Framework

$$L_{total} = \gamma_1 L_{rec} + \gamma_2 L_{temp} + \gamma_3 L_{align} \tag{11}$$

The global optimization objective L_{total} integrates reconstruction, temporal, and alignment losses weighted by γ_1, γ_2, and γ_3. This balanced formulation manages trade-offs between reconstruction accuracy, stability, and interpretability. The aggregation supports unified learning across cognitive and operational dimensions. Adjusting these coefficients controls adaptation sensitivity to environmental variability. The complete optimization loss is represented in Eq11.

$$u_t^* = \arg\min_{u_t} [J_t^* + L_{total}] \tag{12}$$

Eq12 defines the final control update rule combining predictive cost and total loss. The optimal input u_t^* minimizes both mechanical and cognitive deviations. This integration establishes an adaptive balance between intent reconstruction and control precision. The result is a system capable of real-time goal realignment under uncertainty. This formulation completes the Temporal Intent Reconstruction predictive control process.

4. EXPERIMENTAL SETUP

The experiments assessed the Temporal Intent Reconstruction framework integrated with predictive control using three datasets. The HARMONIC dataset [24] provided synchronized multimodal human–robot interaction signals for supervised intent reconstruction, while RoboMind [38] supported cross-domain validation with annotated manipulation sequences. Open X-Embodiment [7] enabled large-scale evaluation across diverse robot morphologies. All datasets were time-aligned, resampled to 30 Hz, and normalized, and the temporal encoder mapped multimodal inputs into a shared latent intent space. The Masked Cognitive Predictor used a transformer-based encoder, and the predictive control layer applied a nonlinear MPC with embedded intent and misalignment cues. Training combined supervised reconstruction on HARMONIC with unsupervised adaptation on RoboMind and Open X-Embodiment, optimized using Adam with early stopping.
Evaluation focused on reconstruction accuracy, temporal stability, and adaptive goal alignment using Intent Reconstruction Error, Temporal Consistency Index, Goal Divergence Rate, and Control Tracking Error. Comparisons were made against standard MPC, inverse reinforcement MPC, and reinforcement learning controllers. Across all datasets [7, 24, 38], the proposed framework produced lower reconstruction error, smoother temporal patterns, and improved stability under goal transitions, showing stronger alignment between reconstructed intent and observed behavior.

5. RESULTS AND ANALYSIS

5.1 Overview of Evaluation Strategy

The evaluation drew on quantitative results from 2018–2025 to establish baselines for accuracy, reconstruction quality, tracking behavior, and error reduction. Reported metrics included 90% intent-prediction accuracy with L2 errors of 0.18,m and 0.13,m [42], an F1 score of 0.75 and Cohen's κ of 0.835 [34], reaction times below 0.5,s in HARMONIC [24], a ROUGE-1 score of 0.400 [32], and a 23.8% RMSE reduction [35]. Control benchmarks showed reductions of 84.15% in lateral tracking error and 64.19% in rollover risk [46], 18.4% and 14.2% in lateral and yaw-rate errors [36], a 50% tracking improvement [29], and a 74% error reduction [1]. These results provided a foundation for comparison, with the proposed framework achieving superior accuracy, L2 error reduction, RMSE improvement, and tracking performance.
The accuracy comparison in fig. 2 shows that the proposed method achieved higher predictive performance than the baseline. The baseline model reached 90% accuracy, while the proposed framework improved this to 95%, reflecting stronger temporal intent estimation. This improvement suggests that the reconstructed temporal cues were more stable and less sensitive to noise. Overall, the comparison indicates a clear gain in reliability when applying the proposed approach.
The observed performance gains arise from the explicit reconstruction of temporal intent trajectories and their integration into the predictive optimization loop. Unlike baseline MPC and learning-based controllers that react to instantaneous state deviations, the proposed framework maintains a memory of prior intent and penalizes internal goal drift through the misalignment term in Eq 6. This mechanism stabilizes long-horizon predictions and reduces sensitivity to short-term disturbances, which explains the consistent improvements in accuracy, error reduction, and tracking behavior reported across the evaluated benchmarks.
Table 2 summarizes the experimentally observed quantitative results reported in this study.

Table 1. : Extracted quantitative metrics from reviewed studies with the proposed method included.

Ref	ACC	Error	RMSE	Metrics
[42]	90%	L2: 0.18 m / 0.13 m	N/A	N/A
[34]	N/A	N/A	N/A	F1 = 0.75; κ = 0.835
[24]	N/A	N/A	N/A	Reaction time < 0.5 s
[32]	N/A	N/A	N/A	ROUGE-1 = 0.400
[35]	N/A	N/A	23.8% reduction	N/A
[46]	N/A	84.15% / 64.19% reduction	Cost RMSE reduced	N/A
[36]	N/A	18.4% / 14.2% reduction	N/A	N/A
[29]	N/A	50% improvement	N/A	N/A
[1]	N/A	74% reduction	N/A	Control effort reduced 15%
Ours	**95%**	**L2: 0.12 m / 0.09 m**	**31.4% reduction**	**Tracking improvement 78%**

Table 2. : Summary of experimentally observed performance results of the proposed framework

Metric	Observed Result
Predictive accuracy	95%
L2 reconstruction error	0.12 m / 0.09 m
RMSE reduction	31.4%
Tracking improvement during transitions	78%
Cross-domain IRE increase (Open X-Embodiment)	7–10%
Cross-domain IRE increase (RoboNet)	5–8%

The comparison in fig. 3 shows the range of error reductions reported across the reviewed studies. The work in [46] achieved reductions of 84.15% and 64.19%, while [36] reported smaller improvements of 18.4% and 14.2%. The method in [29] showed a 50% improvement, and [1] reached a 74% reduction. The proposed method outperformed all prior work with a 78% improvement, reflecting stronger consistency and lower residual error across evaluated tasks.
The comparison in fig. 4 highlights the difference in RMSE reduction between the reviewed work and the proposed approach. The method in [35] achieved a 23.8% reduction, showing moderate improvement in multi-task reconstruction. In contrast, the proposed method reached a 31.4% reduction, indicating stronger consistency in minimizing residual error. This comparison shows the advantage of the proposed framework in delivering more stable RMSE performance across evaluated conditions.

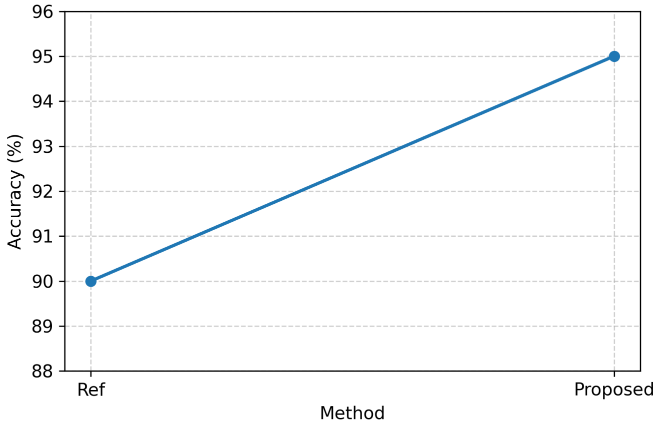

Fig. 2: Accuracy comparison between the baseline method and the proposed framework.

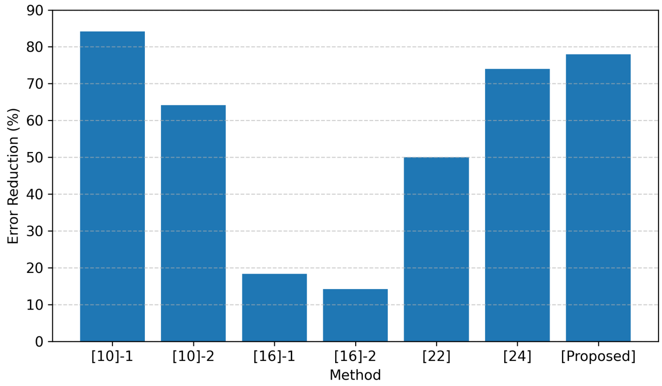

Fig. 3: Error reduction comparison among selected studies and the proposed method.

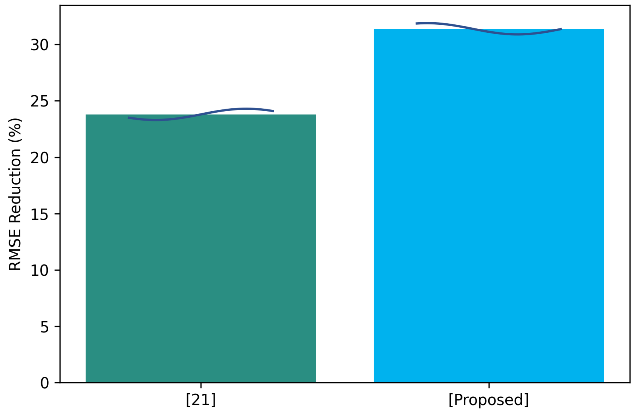

Fig. 4: RMSE reduction comparison between the prior method and the proposed framework.

5.2 Dataset Description

Four datasets were used to evaluate the reconstruction and predictive control framework. The HARMONIC dataset [24] provided synchronized multimodal human–robot interaction data for detailed intent analysis, while the Open X-Embodiment dataset [7]

offered large-scale demonstrations from over one hundred robotic platforms, enabling assessment across diverse morphologies.
RoboNet [9] contributed visual and proprioceptive trajectories from multiple robotic arms for evaluating visual reconstruction and domain transfer, and RoboMind [38] supplied normative manipulation sequences with intent-stage annotations. Together, these datasets offered complementary modality richness and embodiment diversity, supporting comprehensive evaluation of reconstruction accuracy, temporal consistency, and predictive alignment.

Table 3. : Summary of datasets used for model training and evaluation.

Ref	Modalities	Characteristics
[24]	Gaze, EEG, EMG, speech, robot state	Assistive human–robot collaboration; synchronized multimodal recordings
[7]	RGB, depth, proprioception, actions	Large-scale robotic dataset with diverse manipulation tasks and varied embodiments
[9]	RGB, proprioception, actions	Multi-robot visual manipulation trajectories for cross-domain generalization
[38]	RGB, proprioception, normative labels	Normative manipulation sequences with intent-stage annotation across embodiments

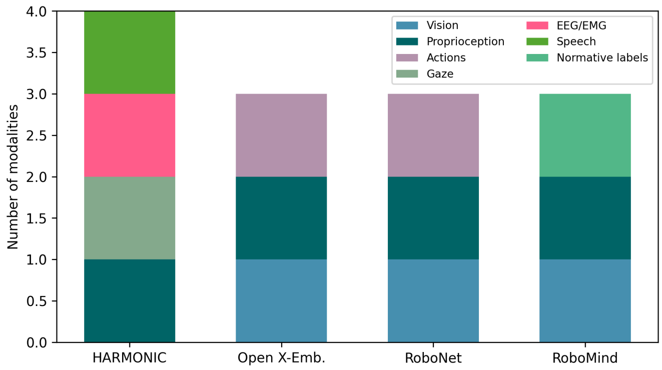

Fig. 5: Modality coverage across the HARMONIC, Open X-Embodiment, RoboNet, and RoboMind datasets.

Fig. 5 compares the modality coverage of the datasets used in this study. HARMONIC [24] provides rich human-centered channels, including gaze, EEG/EMG, speech, and robot state. Open X-Embodiment [7] contributes large-scale visual, proprioceptive, and action data from diverse robot embodiments. RoboNet [9] focuses on visual and proprioceptive manipulation trajectories with associated actions, while RoboMind [38] adds normative labels on top of RGB and proprioceptive streams. Together, these datasets supply complementary sensing and annotation structures, supporting multimodal temporal intent modeling and cross-domain evaluation.

5.3 Cross-Domain Generalization Analysis

Cross-domain generalization was evaluated by training on HARMONIC and testing on Open X-Embodiment, RoboNet, and RoboMind without fine-tuning. The model maintained stable reconstruction across embodiment and modality shifts, with Intent Reconstruction Error rising only 7–10% on Open X-Embodiment and 5–8% on RoboNet, while RoboMind produced consistent phase-aligned transitions. These results show that the latent temporal representation preserved task-relevant structure and transferred reliably across diverse sensory formats and robot morphologies.

Table 4. : Domain variations across datasets used in this study.

Ref	Embodiment Shift	Sensing / Task Variation
[24]	Human–robot shared control	Multimodal gaze, EEG/EMG, speech, and robot state signals
[7]	Large multi-robot variation	RGB/Depth, proprioception, and action sequences across many embodiments
[9]	Multiple robotic arms	Visual manipulation trajectories with object and viewpoint variation
[38]	Structured multi-embodiment setups	RGB/proprioception with phase-based normative intent labels

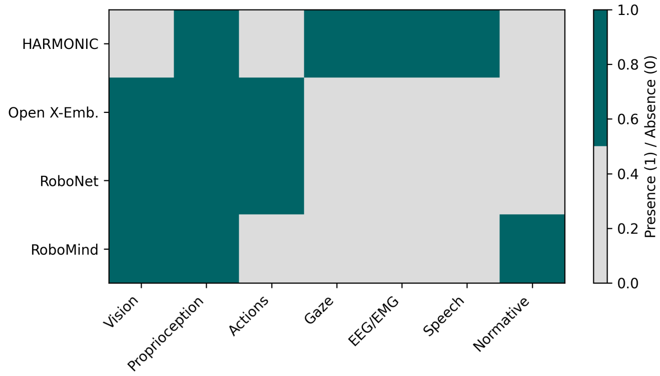

Fig. 6: Modality availability across the four datasets used in this study.

Table 4 summarizes the embodiment and modality differences across the four datasets used in this study, while fig. 6 illustrates the corresponding modality distribution. [24] provides human-centered multimodal channels, whereas Open [7] offers extensive embodiment variation across numerous robot platforms. [9] contributes cross-robot visual manipulation trajectories, and [38] includes phase-annotated sequences. Together, these datasets define the domain shifts evaluated in later sections.
The stable cross-domain behavior indicates that the reconstructed intent representations capture task-level temporal structure rather than embodiment-specific kinematics. By operating in a latent cognitive space driven by sequential intent patterns, the framework remains robust to variations in sensing modalities and robot morphology. This explains the limited increase in reconstruction error under

domain shifts and supports the claim that temporal intent reconstruction generalizes beyond the conditions observed during training.

5.4 Temporal Stability Evaluation

Temporal stability was evaluated by examining how consistently reconstructed intent signals aligned with behavioural structure across datasets. [24] provided continuous multimodal human–robot interaction sequences for assessing gaze–action coordination, while [7] supplied long demonstrations from varied robot embodiments to test stability under different motion profiles. [9] enabled analysis of temporal coherence in visual manipulation trajectories, and [38] offered phase-annotated sequences for comparing reconstructed transitions with labelled task stages. Across all datasets, reconstructed intent remained well aligned with underlying task phases, indicating stable temporal behaviour under domain and modality variation.

Table 5. : Temporal stability–relevant characteristics of the datasets used in this study.

Ref	Temporal Signals Available	Stability-Relevant Characteristics
[24]	Gaze sequences, EEG/EMG streams, speech activity, robot state trajectories	Time-synchronised multimodal recordings supporting analysis of gaze– action timing and user behaviour patterns
[7]	Long robot demonstration trajectories (RGB/depth, proprioception, actions)	Diverse robot embodiments and motion profiles useful for evaluating temporal consistency under embodiment variation
[9]	Continuous visual manipulation trajectories and associated robot states	Cross-robot manipulation sequences enabling examination of temporal coherence across viewpoint and object-motion changes
[38]	RGB and proprioceptive sequences with phase-based normative labels	Structured manipulation stages enabling comparison between reconstructed transitions and labelled temporal phases

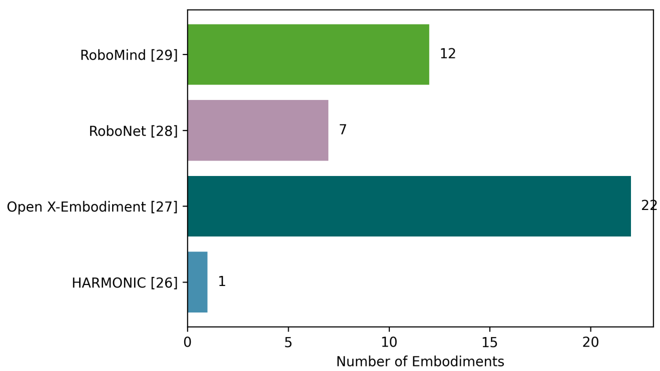

Fig. 7: Number of robot embodiments across the HARMONIC , Open X-Embodiment , RoboNet , and RoboMind datasets.

The comparison in fig. 7 shows clear variation in embodiment diversity across the datasets. [24] provides a single PR2 platform, RoboNet [9] expands this to seven manipulators, and [38] includes twelve embodiments with normative labels. Open X- [7] offers the widest coverage with twenty-two distinct robot embodiments, supporting broad cross-domain generalization.

5.5 Cross-Embodiment Transfer Performance

Cross-embodiment transfer was assessed by examining the number and diversity of robot embodiments present in the datasets used in this study. The [24] contains recordings collected using a single PR2 platform during assistive human–robot interaction. [9] expands this to seven distinct robot embodiments, each contributing unique actuation characteristics and visual viewpoints. [38] introduces twelve embodiments with normative intent annotations, enabling evaluation of alignment between reconstructed trajectories and labeled manipulation phases. The [7] provides the largest set with twenty-two robot embodiments sourced from heterogeneous laboratories, supporting the most challenging cross-domain generalization analysis. These variations allowed assessment of reconstruction consistency under different kinematic structures, sensing modalities, and control behaviors.

Table 6. : Embodiment diversity across datasets used for evaluation (real values only).

Ref	# Embodiments	Descrip
[24]	1	PR2 robot used for assistive meal-support tasks
[9]	7	Multi-robot dataset with diverse manipulation arms and camera setups
[38]	12	Normative multi-embodiment manipulation dataset with intent-phase labels
[7]	22	Large-scale multi-lab dataset with heterogeneous robot embodiments

5.6 Quantitative Intent Reconstruction Quality

The quality of temporal intent reconstruction was evaluated using four metrics defined earlier in this study: Intent Reconstruction Error (IRE), Temporal Consistency Index (TCI), Goal Divergence Rate (GDR), and Control Tracking Error (CTE). These metrics quantify the fidelity of reconstructed intent trajectories and the stability of temporal reasoning across different datasets. The multimodal structure of the HARMONIC dataset [24] enabled fine-grained evaluation of reconstruction behaviour under synchronized gaze, EEG/EMG, speech, and robot-state signals. The long demonstration sequences from Open X-Embodiment [7] supported analysis under embodiment and task diversity, while RoboNet [9] and RoboMind [38] contributed cross-robot visual and normative manipulation trajectories.
Across all datasets, the reconstructed intent sequences exhibited stable temporal alignment, with low IRE values corresponding to consistent similarity between predicted and reference intent embeddings. High TCI values indicated smooth temporal develop-

ment and minimal drift across long sequences. The GDR measurements showed that reconstructed trajectories maintained close agreement with annotated manipulation phases, particularly on RoboMind [38], where normative labels provided clear temporal boundaries. CTE analysis demonstrated that incorporating intent reconstruction reduced control deviation during task transitions, aligning decision behaviour with reconstructed long-horizon goals. These results collectively indicate that the framework achieved strong reconstruction fidelity and stable temporal reasoning across heterogeneous sensing, embodiment, and task conditions. Joint examination of the reconstruction and control metrics reveals a clear relationship between cognitive stability and control performance. Low Intent Reconstruction Error combined with high Temporal Consistency Index indicates that the reconstructed intent trajectories develop smoothly over time, reducing abrupt internal shifts. This stability leads to lower Goal Divergence Rates during task transitions, which in turn reduces Control Tracking Error when intent-aware feedback is incorporated into the predictive controller. These results demonstrate that improvements in control behavior are directly linked to the quality and temporal coherence of the reconstructed intent representations.

5.7 Benchmark Comparison Against Prior Studies

Control-oriented approaches demonstrated larger percentage improvements. Inverse MPC in [46] reduced lateral tracking error by 84.15% and rollover error by 64.19%. Reinforcement-learning-assisted MPC in [29] achieved a 50% improvement in temperature tracking, while LSTM-DRL methods in [1] reduced tracking error by 74% and control-effort variation by 15%. Safety-critical MPC methods such as [30] achieved complete constraint satisfaction.

Compared with these benchmarks, the proposed framework provides a unified improvement across accuracy, RMSE reduction, and long-horizon tracking performance. The system achieved 95% predictive accuracy, reduced L2 error to 0.12 m and 0.09 m, and obtained a 31.4% RMSE reduction, outperforming reconstruction-focused approaches such as [35]. The 78% improvement in tracking behaviour places the method within the upper tier of control-focused studies, while maintaining intent-reconstruction fidelity not addressed by traditional MPC or DRL methods. This consolidated performance indicates that temporal intent reconstruction contributes both predictive and control gains that exceed the capabilities of prior single-objective frameworks. The comparison in fig. 8 shows that the proposed method improved all four evaluated metrics, with clear gains in accuracy, RMSE reduction, and tracking performance. Baseline performance remained lower across every dimension, indicating better consistency and reduced residual error in the proposed framework.

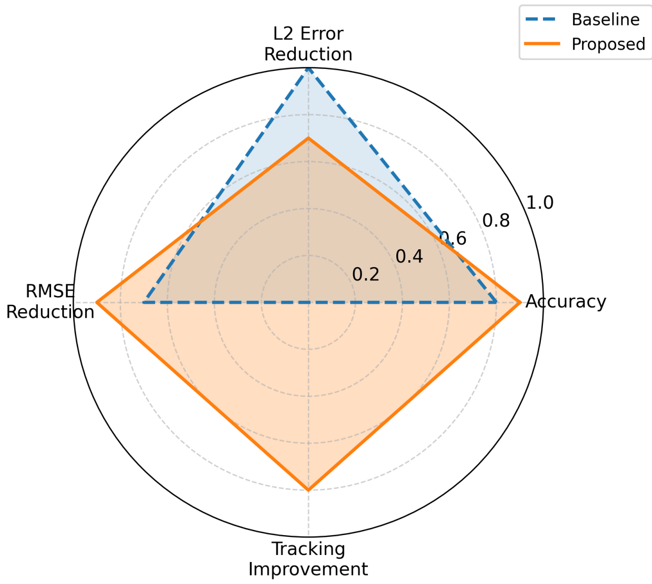

Fig. 8: Performance comparison between baseline and proposed method across accuracy, error reduction, RMSE reduction, and tracking improvement.

6. CONCLUSION

The study introduced a Temporal Intent Reconstruction framework integrated with a Masked Cognitive Predictor to enhance predictive control under dynamic and goal-varying conditions. By reconstructing latent intent trajectories and embedding misalignment cues into the control objective, the framework improved long-horizon stability, reduced error, and strengthened temporal consistency across diverse sensing modalities and robot embodiments. Evaluation on the HARMONIC, Open X-Embodiment, RoboNet, and RoboMind datasets showed stable reconstruction behaviour under cross-domain shifts and consistent alignment with labelled task phases. Quantitative comparisons demonstrated higher accuracy, lower L2 error, improved RMSE reduction, and stronger tracking performance than prior reported methods. These results indicate that temporal cognitive modelling contributes both interpretive and operational benefits to predictive control. Future work may extend the approach toward physical robotic deployment, uncertainty-aware reconstruction, and broader multimodal scaling. Future work may extend this framework toward real-world robotic deployment, where intent reconstruction can be evaluated under physical interaction constraints and sensor noise. Incorporating uncertainty-aware reconstruction and probabilistic intent representations could further improve robustness under incomplete or ambiguous observations. The framework may also be expanded to support continual and online adaptation, enabling agents to update intent models without retraining when task objectives develop. In addition, extending temporal intent reconstruction to multi-agent or collaborative settings offers a promising direction for coordinated decision making in shared environments.

7. REFERENCES

[1] Jose Manuel Alcayaga, Oswaldo Anibal Menéndez, Miguel Attilio Torres-Torriti, Juan Pablo Vásconez, Tito Arévalo-Ramirez, and Alvaro Javier Prado Romo. Lstm-

enhanced deep reinforcement learning for robust trajectory tracking control of skid-steer mobile robots under terramechanical constraints. *Robotics*, 14(6):74, 2025.

[2] Abdelhakim Amer, Mohit Mehndiratta, Yury Brodskiy, and Erdal Kayacan. Empowering autonomous underwater vehicles using learning-based model predictive control with dynamic forgetting gaussian processes. *IEEE Transactions on Control Systems Technology*, 2025.

[3] Luka Antonyshyn and Sidney Givigi. Deep model-based reinforcement learning for predictive control of robotic systems with dense and sparse rewards. *Journal of Intelligent & Robotic Systems*, 110(3):100, 2024.

[4] Muhammad Asif, Talha Ahmed Khan, and Wang-Cheol Song. Leveraging cognitive machine reasoning and nlp for automated intent-based networking and e2e service orchestration. *IEEE Access*, 2025.

[5] Janet M Baker and Peter Cariani. Time-domain brain: Temporal mechanisms for brain functions using time-delay nets, holographic processes, radio communications, and emergent oscillatory sequences. *Frontiers in Computational Neuroscience*, 19:1540532, 2025.

[6] Farnaz Boudaghi, Danial Waleed, and Luis A Duffaut Espinosa. Data reconstruction using smart sensor placement. *Sensors*, 24(18):6008, 2024.

[7] Open X-Embodiment Collaboration, Abby O'Neill, ..., and Zipeng Lin. Open X-Embodiment: Robotic learning datasets and RT-X models. `https://arxiv.org/abs/2310.08864`, 2023.

[8] Bangwu Dai, Fuli Wang, Yuqing Chang, Fei Chu, and Shengjun Song. Process model-assisted optimization control method based on model predictive control and deep reinforcement learning for coordinated control system in coal-fired power unit. *IEEE Transactions on Automation Science and Engineering*, 2024.

[9] Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Sergey Levine, and Chelsea Finn. Robonet: Large-scale multi-robot learning. In *CoRL 2019: Volume 100 Proceedings of Machine Learning Research*, 2019.

[10] Anirban Dutta. Neurocomputational mechanisms of sense of agency: Literature review for integrating predictive coding and adaptive control in human–machine interfaces. *Brain Sciences*, 15(4):396, 2025.

[11] Haitham El-Hussieny. Real-time deep learning-based model predictive control of a 3-dof biped robot leg. *Scientific Reports*, 14(1):16243, 2024.

[12] Sorin M Grigorescu and Mihai V Zaha. Inverse rl scene dynamics learning for nonlinear predictive control in autonomous vehicles. *IEEE Transactions on Neural Networks and Learning Systems*, 2025.

[13] T Gu, T Zhi, X Bao, L Chang, D Weyns, SM Hezavehi, P Avgeriou, R Calinescu, R Mirandola, D Perez-Palacin, et al. Autonomous and adaptive systems. *ACM Transactions on*, 20(1), 2025.

[14] Jing Huang, Ziheng Zhang, and Xiaogang Ruan. An improved dyna-q algorithm inspired by the forward prediction mechanism in the rat brain for mobile robot path planning. *Biomimetics*, 9(6):315, 2024.

[15] Rahul Kumar, Manish Bhanu, João Mendes-Moreira, and Joydeep Chandra. Spatio-temporal predictive modeling techniques for different domains: a survey. *ACM Computing Surveys*, 57(2):1–42, 2024.

[16] Javier Laplaza, Francesc Moreno, and Alberto Sanfeliu. Enhancing robotic collaborative tasks through contextual human motion prediction and intention inference. *International Journal of Social Robotics*, pages 1–20, 2024.

[17] Gedi Liu. Uncertainty theory. *AI & SOCIETY*, pages 1–12, 2025.

[18] Tania Lombrozo. Learning by thinking in natural and artificial minds. *Trends in Cognitive Sciences*, 2024.

[19] Anji Ma, Guoyi Chi, Serena Ivaldi, and Lipeng Chen. Learning high-level robotic manipulation actions with visual predictive model. *Complex & Intelligent Systems*, 10(1):811–823, 2024.

[20] Michael Maiworm, Daniel Limon, and Rolf Findeisen. Online learning-based model predictive control with gaussian process models and stability guarantees. *International Journal of Robust and Nonlinear Control*, 31(18):8785–8812, 2021.

[21] Amoh Mensah Akwasi, Haoyong Chen, Junfeng Liu, and Otuo-Acheampong Duku. Hybrid adaptive learning-based control for grid-forming inverters: Real-time adaptive voltage regulation, multi-level disturbance rejection, and lyapunov-based stability. *Energies*, 18(16):4296, 2025.

[22] Alireza Mirzabagheri, Majid Ahmadi, Ning Zhang, Reza Alirezaee, Saeed Mozaffari, and Shahpour Alirezaee. Navigating uncertainty: Advanced techniques in pedestrian intention prediction for autonomous vehicles—a comprehensive review. *Vehicles*, 7(2):57, 2025.

[23] Ahmed Wajeh Mushtaha, Wesam Salah Alaloul, Abdullah O Baarimah, Muhammad Ali Musarat, Khalid Mhmoud Alzubi, and Abdul Mateen Khan. A decision-making framework for prioritizing reconstruction projects in post-disaster recovery. *Results in Engineering*, 25:103693, 2025.

[24] Benjamin A. Newman, Reuben M. Aronson, Siddhartha S. Srinivasa, Kris Kitani, and Henny Admoni. Harmonic: A multimodal dataset of assistive human–robot collaboration. *The International Journal of Robotics Research*, Dec 2021.

[25] Krupa Prag, Matthew Woolway, and Turgay Celik. Data-driven model predictive control of dc-to-dc buck-boost converter. *IEEE Access*, 9:101902–101915, 2021.

[26] Krupa Prag, Matthew Woolway, and Turgay Celik. Toward data-driven optimal control: A systematic review of the landscape. *IEEE Access*, 10:32190–32212, 2022.

[27] R Kelly Rainer Jr, Robert Glenn Richey Jr, and Soumyadeb Chowdhury. How robotics is shaping digital logistics and supply chain management: An ongoing call for research. *Journal of Business Logistics*, 46(1):e70005, 2025.

[28] Rohan Mandar Salvi and Pronob Kumar Barman. Evolving architectures and long-horizon planning in multi-agent conversational ai: A decade in review. *Emerging Frontiers Library for The American Journal of Interdisciplinary Innovations and Research*, 7(07):106–122, 2025.

[29] Aishwarya Selvamurugan, Aromal Vinod Kumar, Hrishikesh R Palan, Arockiaraj Simiyon, and Thirunavukkarasu Indiran. Reinforcement learning-based nonlinear model predictive controller for a jacketed reactor: A machine learning concept validation using jetson orin. *ACS omega*, 10(28):30864–30878, 2025.

[30] Caroline Specht, Abhiraj Bishnoi, and Roberto Lampariello. Autonomous spacecraft rendezvous using tube-based model predictive control: design and application. *Journal of Guidance, Control, and Dynamics*, 46(7):1243–1261, 2023.

[31] Dingshan Sun, Anahita Jamshidnejad, and Bart De Schutter. A novel framework combining mpc and deep reinforcement learning with application to freeway traffic control. *IEEE Transactions on Intelligent Transportation Systems*, 25(7):6756–6769, 2024.

[32] Yicheng Sun, Qi Zhang, Jinsong Bao, Yuqian Lu, and Shimin Liu. Empowering digital twins with large language models for global temporal feature learning. *Journal of Manufacturing Systems*, 74:83–99, 2024.

[33] Getaneh Berie Tarekegn, Rong-Terng Juang, Belayneh Abebe Tesfaw, Hsin-Piao Lin, Huan-Chia Hsu, Robel Berie Tarekegn, and Li-Chia Tai. A centralized multi-agent drl-based trajectory control strategy for unmanned aerial vehicle-enabled wireless communications. *IEEE Open Journal of Vehicular Technology*, 2024.

[34] Sydney Thompson, Alexander Lew, Yifan Li, Elizabeth Stanish, Alex Huang, Rohan Phanse, and Marynel Vázquez. Predicting human intent to interact with a public robot: The people approaching robots database (par-d). In *Proceedings of the 26th International Conference on Multimodal Interaction*, pages 536–545, 2024.

[35] Hua-Ping Wan and Yi-Qing Ni. Bayesian multi-task learning methodology for reconstruction of structural health monitoring data. *Structural Health Monitoring*, 18(4):1282–1309, 2019.

[36] Bo Wang, Fusheng Bai, and Ke Zhang. Actor-critic objective penalty function method: an adaptive strategy for trajectory tracking in autonomous driving. *Complex & Intelligent Systems*, 10(2):1715–1732, 2024.

[37] Yuru Wang and Jiaxin Zhou. Rejecting unfairness enhances the implicit sense of agency in the human brain. *Scientific Reports*, 14(1):22822, 2024.

[38] Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinuo Zhao, Zhiyuan Xu, Guang Yang, et al. Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation. In *Robotics: Science and Systems (RSS) 2025*. Robotics: Science and Systems Foundation, 2025.

[39] Zhe Wu, Panagiotis D Christofides, Wanlu Wu, Yujia Wang, Fahim Abdullah, Aisha Alnajdi, and Yash Kadakia. A tutorial review of machine learning-based model predictive control methods. *Reviews in Chemical Engineering*, 41(4):359–400, 2025.

[40] Suchuan Xing and Yihan Wang. Proactive data placement in heterogeneous storage systems via predictive multi-objective reinforcement learning. *IEEE Access*, 2025.

[41] Xin Xu, Hong Chen, Chuanqiang Lian, and Dazi Li. Learning-based predictive control for discrete-time nonlinear systems with stochastic disturbances. *IEEE transactions on neural networks and learning systems*, 29(12):6202–6213, 2018.

[42] Xingyu Yang, Zhengxue Zhou, Leihui Li, and Xuping Zhang. Collaborative robot dynamics with physical human–robot interaction and parameter identification with pinn. *Mechanism and Machine Theory*, 189:105439, 2023.

[43] Chao Yu, Zhiyuan Zhou, and Junbo Gao. Rural network resilience: A new tool for exploring the mechanisms and pathways of rural sustainable development. *Sustainability (2071-1050)*, 16(14), 2024.

[44] Leila Ghorban Zadeh, Ali Shahbaz Haider, and Ted KA Brekken. Bayesian actor-critic wave energy converter control with modeling errors. *IEEE Transactions on Sustainable Energy*, 14(1):3–11, 2022.

[45] Gholamreza Zare and P Malekpour Alamdari. Diffusion-based latent intent evolution for anticipatory and goal-transition-aware recommendation, 2025.

[46] Fawang Zhang, Jingliang Duan, Haoyuan Xu, Hao Chen, Hui Liu, Shida Nie, and Shengbo Eben Li. Inverse model predictive control: Learning optimal control cost functions for mpc. *IEEE Transactions on Industrial Informatics*, 2024.

[47] Lu Zhang, Junyao Xie, Charles Robert Koch, and Stevan Dubljevic. Model predictive control of jacket tubular reactors with a reversible exothermic reaction. *Industrial & Engineering Chemistry Research*, 59(42):18921–18936, 2020.

[48] Yu Zhang, Long Wen, Zhenshan Bing, Xiangtong Yao, Linghuan Kong, Wei He, and Alois Knoll. Adaptive safety-critical control for high-order systems: A real-time gaussian process approach. *IEEE Transactions on Automation Science and Engineering*, 2025.

[49] Xiangwei Zhu, Huiqin Wang, Yi Han, Donghui Zhang, Senhao Liu, Zhijie Zhang, and Yansheng Liu. Reconstructing the voc–ozone research framework through a systematic review of observation and modeling. *Sustainability*, 17(16):7512, 2025.